

Landmark Based Position Estimation for Pinpoint Landing on Mars

Yang Cheng and Adnan Ansar
Mobility Systems Concept Development Section
Jet Propulsion Laboratory, California Institute of Technology
4800 Oak Grove Drive, Pasadena, CA 91109-8099, USA
ycheng@telerobotics.jpl.nasa.gov

Abstract - Many of NASA's planned missions for the coming decade will require a pinpoint landing (PPL) capability, whether for sample acquisition and return or for precise insertion into hazardous but scientifically interesting terrain. Thus, a robust spacecraft-based position estimation system is a critical near term need. We present in this paper a vision-based system, which provides a low power, low cost, high accuracy solution with flight proven hardware. The key components of our image based localization approach are: (1) real-time detection of landmarks in descent imagery, (2) real-time matching of detected landmarks to a stored database, and (3) robust estimation of spacecraft state (position, attitude and velocity) from matched 2D-3D data. We present an analysis of the performance and noise sensitivity of our system and show that the suggested technology is able to deliver a spacecraft to within 100 meters of a pre-selected landing site in a typical Mars landing scenario.

Index Terms: Pinpoint landing, landmark detection, landmark matching, position estimation

I. INTRODUCTION

Current descent and landing technology for planetary missions, such as those to Mars, is characterized by a greater than 30 x 100 km landing error ellipse with no terrain recognition or hazard avoidance mechanism. In the next decade, NASA plans to demonstrate an ambitious capability – pinpoint landing on another planetary body. The primary objective of PPL is to deliver a spacecraft to within 100 meters of a targeted landing site. In order to accomplish this, several new technologies will be developed over the next few years. These include new optical guided navigation systems for precision entry, an optimized powered descent guidance system, advanced parachute technology, a new propulsive traverse capability, and a new real-time terrain recognition capability for spacecraft localization.

NASA has successfully flown automated landing systems for Lunar and Mars missions, such as *Surveyor*, *Apollo*, *Viking*, *Mars Pathfinder* and *Mars Exploration Rover (MER)*. However, only one of these missions had a terrain-relative guidance, navigation and control

(GNC) system. In MER, the Descent Image Motion Estimation System (DIMES) was used to estimate horizontal velocities. DIMES consisted of a descent imager, a radar altimeter, an inertial measurement unit and an algorithm to provide a low cost, robust and computationally efficient solution to the horizontal velocity estimation problem, which was critical for the safe deployment of the air bags during touchdown [9]. DIMES is the first ever terrain-relative sensing and guidance system used by a real mission. However, unlike our proposed system, DIMES has neither terrain recognition nor hazard avoidance capabilities.

Our real-time spacecraft localization scheme, the first of its kind, works as follows. A landing site on the targeted body is selected on Earth using orbital imagery, and the landmarks (e.g. craters) within the landing ellipse are mapped. During descent, the lander's initial position with respect to the landmarks as well as to the selected landing site is determined automatically on board. The lander is then guided to the landing site using continuous updates of lander position and velocity throughout the descent. There are three fundamental requirements for this new technology.

1. The system must be able to recognize the terrain (landmarks) reliably and repeatedly over a wide variety of environments. During spacecraft descent, considerable environmental variations, such as lighting angle, atmospheric conditions, viewing angle, and spacecraft altitude, can affect the appearances of landmarks. The key to success relies on defining a class of landmarks that has good invariance properties under variable environmental conditions and a set of associated algorithms for handling this type of landmark robustly.
2. The system must accomplish the task under extreme conditions imposed by the slow flight computer and terminal descent time constraints. The projected CPU clock speed of the flight computer for a PPL mission will be on the order of 100 MIPS, which is not ideal for processing large volumes of data. For instance, a descent image is typically 1 MB in size. There is a roughly 60 second window of opportunity for spacecraft localization during the parachute stage between the heat shield jettison and powered descent. In order to obtain reliable spacecraft state information, several images as well as other

onboard sensor data such from an IMU and altimeter must be processed during this period. In general, each cycle of spacecraft localization, which includes image acquisition, image processing and sensor data infusion, should be done in a few seconds. This is a very demanding requirement.

3. The system must be able to guide the spacecraft to land within 100 meters of the target under noisy conditions. Due to the limitations of both hardware and software, system noise can only be compensated to a limited extent. The primary noise sources are the landmark detection error, the base map error (both position and elevation), sensor noises (imager, IMU, altimeter), and image and IMU sensor misalignment. All of these influence the performance of the system. By considering these uncertainties, the system must be robust enough to provide valid spacecraft state, which meets or exceeds the PPL requirement.

The rest of this paper, we will present a system that can meet these requirements in the case of crater landmarks, which appear densely on many bodies of interest.

II. SYSTEM DESCRIPTION

Craters are landforms commonly found on the surface of planets, satellites, asteroids, and other solar system bodies. A crater, in general, is a bowl shaped depression created by collision or volcanic activity. Because of their simple and unique geometry and relatively stable appearance under different viewing and lighting angles, craters are ideal landmarks for spacecraft localization [1-3]. A large number of craters can be found on the surface of Mars, which can generally be divided into three regions – heavily, moderately and lightly cratered [5, 10]. A statistical study shows that there will be adequate craters in a landing ellipse to ensure positive spacecraft localization. Therefore, we have chosen to focus on craters as our landmark of choice.

Real-Time Crater Detection

A real-time crater detection algorithm has been developed based on previous work in crater detection algorithm for autonomous spacecraft navigation [2]. This algorithm breaks down into five steps:

1. **Edge Detection:** This step detects edges in an image and places them in a database.
2. **Rim Edge Grouping:** This step groups together edges that belong to the same crater. The information used for this process includes edge shape (convex), the image intensity profile inside a crater, and edge gradients. If a pair of edges corresponding to the lit and shaded side of a single crater is found, they will be used to fit an ellipse.
3. **Ellipse Fitting:** This step fits an ellipse to each group of crater edges using an iterative algorithm –

the reweighting least squares method, which robustly removes outliers in the input data.

4. **Precision Fitting:** This step adjusts the detected crater's geometry directly in the image domain to reduce errors introduced in edge detection and ellipse fitting. A multidimensional iterative nonlinear minimization algorithm based on conjugate gradients is used to lock an ellipse precisely onto the rim of a crater.
5. **Crater Confidence Evaluation:** This step evaluates every detected crater and assigns a confidence value to it.

After extensive optimization, the algorithm outlined above is able to meet the PPL time and performance requirements for the first time. The performance improvement stems from the following modifications

1. Reorganization of data to allow maximum usage of pointer operations.
2. Conversion and analysis of image features in vector data format.
3. Use of a hash table vector database for fast data retrieval.
4. Extensive use of lookup tables for repetitive computations.

The current algorithm is able to detect craters from a 512 x 512 image in less than 0.8 second on a 333 MIPS Ultra SUN station, thus meeting the Mars PPL requirement.

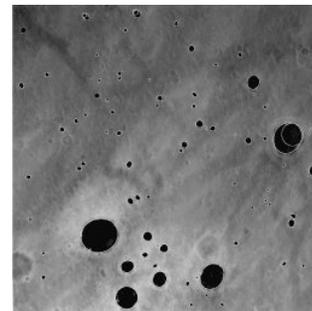


Figure 1: Example of a crater detection result from MER-A descent imagery. The detected craters are shaded for easy visualization.

Extensive experimentation shows that the detection rate is better than 94% and the false alarm rate is less than 7%. The position error is less than 0.3 pixel and the geometric error in shape is less than 0.5 pixel.

Crater Matching

We use geometric recognition techniques to match craters extracted from an image to a database containing the 3D locations of the craters. Each crater is treated as an attributed point corresponding to the center of the crater, where the attributes are the radius and orientation of the crater. The efficiency of the basic method is improved by two means. First, the crater attributes are

used to remove matches that are incompatible. Second, an initial estimate of the spacecraft position is used to filter matches that are not feasible. The information from other sensors such as altitude from the altimeter and attitude from the IMU can help to reduce the search scope even further. The attitude is in the inertial frame and available at all times. The offset between the inertial frame and local map, which could be off by a few degrees, is incorporated into the matching. The altitude will not be available until the lander is a few kilometers (3 km is typical) from the ground. Until then a larger scale range is searched. However the problem can be reduced by use of conic invariants[4]. A pair of coplanar conics c_1 and c_2 has two invariants

$$\begin{aligned} I_{c_1c_2} &= \text{Trace}(c_1^{-1}c_2) \\ I_{c_2c_1} &= \text{Trace}(c_2^{-1}c_1) \end{aligned}$$

Since under a linear transformation $\mathbf{x} = \mathbf{TX}$, c_1 and c_2 map to $C_1 = \mathbf{T}'c_1\mathbf{T}$, and $C_2 = \mathbf{T}'c_2\mathbf{T}$, we have

$$\begin{aligned} I_{C_1C_2} &= \text{Trace}(\mathbf{T}^{-1}c_1^{-1}(\mathbf{T}')^{-1}\mathbf{T}'c_2\mathbf{T}) \\ &= \text{Trace}(c_1^{-1}c_2) = I_{c_1c_2} \end{aligned}$$

The same derivation holds for $I_{c_2c_1}$. These invariant moments are also valid under roll, pitch and shearing because the relationship of a planar surface between two image views is a homography.

In Fig. 2 we show the result of matching the craters detected in Fig. 1 to a 2048 x 8366 Mars Orbiter Camera (MOC) image strip with approximately 1.41 meter/pixel resolution and 1,777 detected craters. The MOC data serves as our map for this test. As shown in Fig. 2, the crater matching algorithm successfully matches the craters from the descent image to the map.

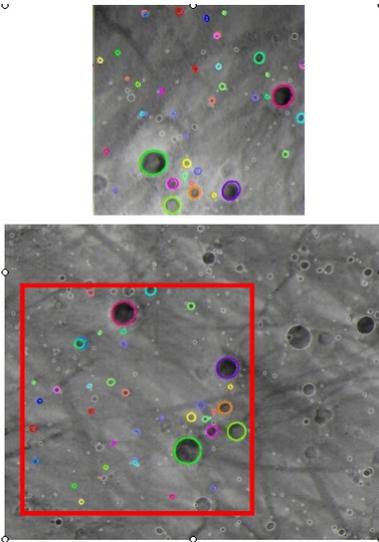


Figure 2: An example of the crater matching result between MER descent image (top) and MOC image. The color rim indicates the correspondences.

Robust Position Estimation

From a suitable number of matched landmarks for which we have prior geometric data, we accurately estimate the position and orientation of the spacecraft with respect to the surface of the planetary body. In the case of craters, the relevant structures are crater centroids, which are estimated carefully to account for perspective distortion effects.

Given a collection of points in 3D and their 2D projections, we recover camera pose as follows. The 3D points are originally presented with respect to some reference coordinate frame, typically dependent on the landing ellipse and independent of the location of the camera. The first step in recovering camera pose is to determine the coordinates of these points in a coordinate frame centered on the camera. From prior calibration of the camera, we know the exact 2D coordinates of a pixel on the image plane (CCD or CMOS device). If (x,y) are the 2D coordinates of an image point p arising from a 3D point P , then P can be expressed in the coordinate frame of the camera as (I_x, I_y, I) for some suitable scale factor I . Note that the distances between 3D points are independent of coordinate system. Hence, for a collection of image points $\{p_i\}$ and associated 3D points $\{P_i\}$, we know $\{d_{ij} = \|P_i - P_j\|\}$. This can be expressed as

$$d_{ij}^2 = (I_i x_i - I_j x_j)^2 + (I_i y_i - I_j y_j)^2 + (I_i z_i - I_j z_j)^2$$

resulting in a set of quadratic equations in the unknown $\{I_i\}$. We use an efficient and robust linear algorithm to solve for the I_i [6]. Once these quantities are known, the 3D coordinates of all points are known in both the coordinate frame of the camera and the reference frame. Recovery of the camera pose is then equivalent to finding the Euclidean transformation that maps one of these point clouds onto the other. This absolute orientation problem has several known linear solutions [6]. For small numbers of points, the proposed method is fast and robust. It requires no initialization because there is no iterative component. Consequently, there is no issue with slow convergence or local minima. Given an initial estimate from the above algorithm, we use a fast iterative method [8] to refine the result. Finally, we use a robust estimation approach to decrease sensitivity to outliers. We perform position estimation over a statistically meaningful number of trials using a subset of the data and a Least Median Squares (LMedS) criterion based on image reprojection errors.

III. PERFORMANCE ANALYSIS

The ultimate objective of pinpoint landing is to deliver a spacecraft to within 100 meters of a targeted site. However, the uncertainties of hardware, software and landing site topography can influence system performance. The primary noise sources are landmark detection error, base map error (both position and elevation), sensor noise (imager, IMU, altimeter), and image and IMU sensor misalignment. Although these

uncertainties cannot be eliminated, we show that their influence will not drive the estimate beyond the required landing error envelope. In addition, terrain recognition, (i.e. crater matching) must be unambiguous. In this section, we study the influence of a few major uncertainties.

Crater Constellation Uniqueness Analysis

To be used for unambiguous position estimation, the constellation of landmarks must be unique in terms of size and location in the landing ellipse. We have conducted two studies in this area. The first uses an Odyssey THEMIS image containing 917 detected craters (Fig. 3). A call a set of neighboring craters a crater constellation. We performs a pairwise comparison of constellations with regard to size and relative configuration (to within 1 pixel) to determine the likelihood of ambiguity in the selection.

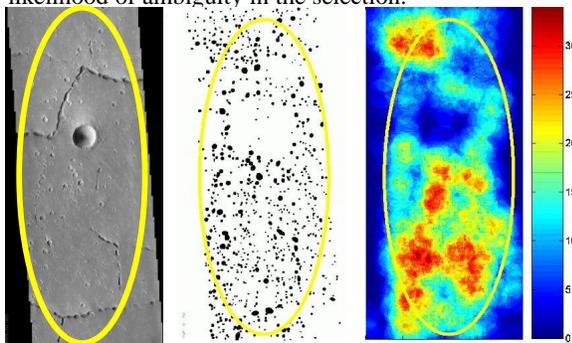


Figure 3: Odyssey THEMIS image (left). Detected craters (center). Crater density map (right)

The probability of an ambiguous constellation in this dataset is shown in Table 1, which indicates that given the size and position of craters, the probability of an ambiguous configurations is very small ($< 0.0001\%$) when the number of craters in a constellation is greater than 5.

Table 1: Probability of ambiguous configuration from a real crater database

Number of craters	Total # valid configurations	# Ambiguous configurations	% probability of confusion
2	7.1×10^5	71,898	10.1
3	1.3×10^7	7,394	5.8×10^{-2}
4	1.8×10^8	1027	5.7×10^{-4}
5	2.2×10^9	221	1.0×10^{-5}

The second study is a statistical analysis of the likelihood of confusion given a model for crater distribution [10]. This again takes into account the crater sizes and relative distances. We present a sketch of the derivation and the results. The probability of having two pairs of craters at the same relative distance d in a disk (region of attention) of radius R can be computed as:

$$P_1 = \left(\frac{3}{4} \frac{1}{R} \arctan \left(\sqrt{\frac{4R^2}{d^2} - 1} \right) \right) \left(\left(\frac{d}{4} \right)^2 \frac{3}{R^3} \sqrt{\frac{4R^2}{d^2} - 1} \right)$$

We omit the complete derivation of this and some other quantities for space considerations.

Assuming a tolerance d_p for distances to be considered "equal," the probability of having any two pairs of craters at the same relative distance in the disk is given by integrating P_1 over all admissible distances to get a probability P_2 . Suppose there are N total craters in the disk and we use m of them for position estimation, leaving $n = N - m$. The probability of having another m craters with the same geometric configuration can be computed as

$$P_{geom} = \left(1 - \left(1 - \frac{d_p^2}{R^2} \right)^n \right)^{m-2} \left(1 - (1 - P_2)^{(2n-3)(2m-3)} \right)$$

Given a model for crater distribution from [10], we find by integrating over all craters from size d_{min} to d_{max} that the total number of craters in our disk is:

$$N = A \cdot \frac{K}{1-a} \left(d_{max}^{1-a} - d_{min}^{1-a} \right)$$

where A is the area of the disk and K and a determine crater size distribution. This also gives a probability density $s_N(d_{min}, d_{max}) = N/A$ for craters within the given size. The probability of having two craters with the same size (up to a tolerance d_s) can be computed by integrating $s_N(d, d+d_s)^2$ over all admissible d . Let this quantity be P_3 . Then the probability of m craters out of a candidate population of n having the same size as our m selected craters is

$$P_{size} = \left(1 - (1 - p_3)^n \right)^m$$

Finally, the probability of an ambiguous constellation is the product of the size probability and the geometric configuration probability

$$P_{tot} = P_{size} P_{geom}$$

For a heavily cratered region, we get $(K, a) = (0.3, 1.8)$ from [10]. With disk diameter = 8 km, $d_s = d_p = 30$ m, $d_{max} = 4$ km and $d_{min} = 100$ m, we obtain the results in Table 2. These match the results in Table 1 up to order of magnitude, which is as much as we can expect for this relatively vague notion of "heavy" cratering..

Table 2: Probability of ambiguous configuration from crater distribution models.

Number of craters	Lightly cratered	Heavily cratered
2	5.4	44.4
3	5.2×10^{-3}	0.20
4	2.9×10^{-6}	8.4×10^{-4}
5	1.4×10^{-9}	3.4×10^{-6}

Position Estimation Analysis

From a suitable number of matched landmarks for which we have prior geometric data, we can accurately estimate the position and orientation of the spacecraft with respect to the surface of the planetary body. We identify 3D to 2D point correspondences between our

database and descent or orbital imagery. In the case of craters, the relevant structures are crater centroids. We study the accuracy of this approach via a detailed set of simulations parametrized by insertion altitude, orientation, imager resolution and camera field of view (FOV). The virtual onboard camera takes a snapshot of the visible portion of the 3D terrain. The true 2D image coordinates of the crater centroids are distorted with zero mean Gaussian noise of a given standard deviation and truncated beyond some maximum value. If there are not enough craters (a tunable parameter n with a minimum value of 4) for position estimation, we mark the insertion point as a failure. Otherwise, we randomly select n of the visible craters for position estimation. Assuming matches between the 2D and 3D datasets, we use our position estimation algorithm to estimate spacecraft position and orientation and compare to the preset ground truth. For simulation purposes, we corrupt image data with noise with standard deviation of 1 pixel. In Fig. 4 we show the result for position estimation for our base algorithm and for the statistically robust version described previously. In each case, we assume 12 point matches and a 1 pixel random image noise with a variable percentage of outliers. The camera is assumed to be 8 km above the surface.

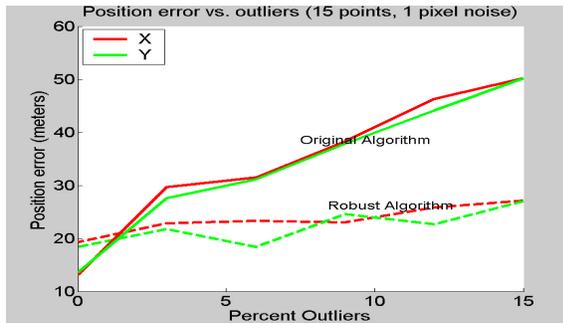


Figure 4: Position estimation error as a function of outliers in matched craters

Spacecraft Velocity Estimation Analysis

As a prelude to integration of a full Kalman filter solution for velocity estimation, we evaluate a simple estimation scheme based on integrated IMU acceleration data and position estimates provided by our algorithms above. Our analysis gives an upper bound for expected velocity error.

Without acceleration data, two camera frames provide an estimate of average velocity only. However, if acceleration is present, we can compute an instantaneous velocity depending on IMU sampling accuracy by integration. If the frames are taken at times t_0 and t_f with recovered positions $P(t_0)$ and $P(t_f)$ we compute $V(t)$ from the acceleration $a(t)$ as follows:

$$V(t) = \frac{P(t_f) - P(t_0) - \int_{t_0}^{t_f} a(t) dt}{t_f - t_0} + \int_{t_0}^t a(t) dt$$

If more than two frames are available, we take all pairs and compute a weighted sum. We have determined empirically that weighting linearly by time interval between frames and inversely by distance from the ground plane at time of frame capture works well. In Fig. 5., we plot the velocity error for a trial trajectory with varying frame numbers and frame intervals. With the current simulation framework in place, we will be able to do much more sophisticated analysis of velocity estimation using varying trajectories and error models for IMU and image noise as well as crater matching.

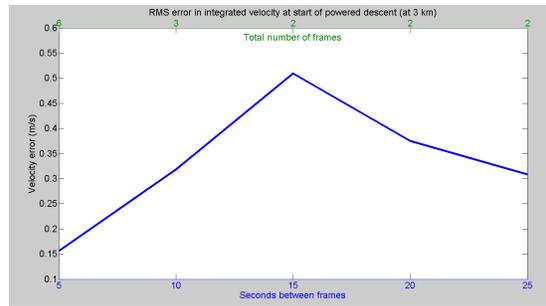


Figure 5: Velocity estimation error using varying frame numbers and intervals.

With the simple scenario tested, we obtain velocity estimates at least an order of magnitude better than those currently available in mission scenarios.

Landing Error Analysis

We consider only the landing error arising from our position estimation system. We do not attempt to duplicate the sophisticated guidance algorithms to be used during a powered descent. Instead we simply report landing error arising from vision based position estimates as the sum of the position estimate at start of powered descent and the drift due to instantaneous velocity estimation error. Thus, if LE = landing error, PE = position estimation error at powered descent, VE = velocity error at powered descent, and t = time remaining to the ground, we have

$$LE = PE + VE \cdot t$$

In Fig. 6 we show the landing error for the case of the trajectory discussed above for various starting points of powered descent. Observe, that for this case of perfect acceleration knowledge, we are well within the requirements for pinpoint landing with integrated IMU data.

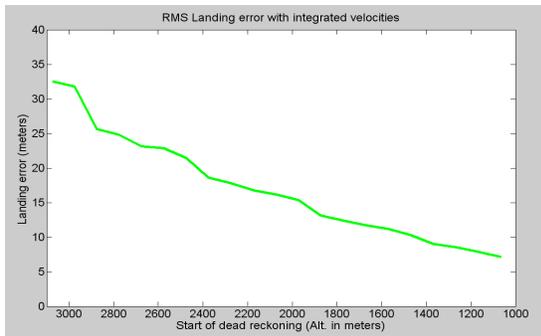


Figure 6: Landing error for MSL-like trajectory assuming IMU integrated (green)

If the base map errors (both position and elevation) are also considered, the landing error increases. Fig. 7 and 8 show the landing error vs. map error for start of powered descent at 1 km and 3 km, respectively. We see that the landing error is well within the requirements of pinpoint landing if the map and elevation error are less than 5 meters.

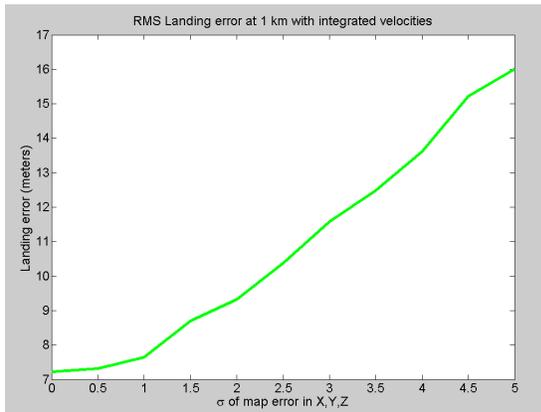


Figure 7: RMS landing error vs. map error when the spacecraft starts dead reckoning at 1 km.

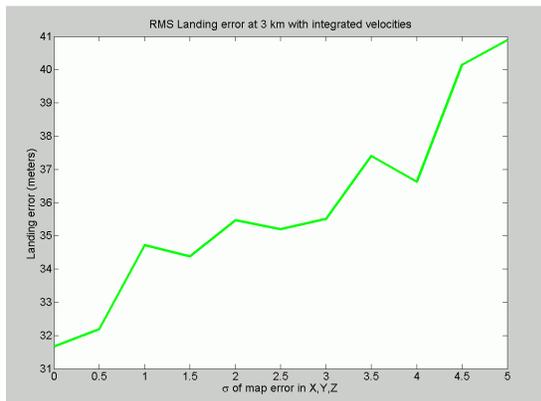


Figure 8: RMS landing error vs. map error when the spacecraft starts dead reckoning at 3 km.

IV. CONCLUSIONS

In this paper, a landmark (crater) based position estimation system is suggested. The performance analysis shows that this system is able to guide a spacecraft to land on Mars within 100 meter of a targeted landing site under likely noise conditions.

Acknowledgments

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

References

- [1] Yang Cheng, A. Johnson, C. Olson, L. Matthies, "Optical Landmark Detection for Spacecraft Navigation" 13th Annual AAS/AIAA Space Flight Mechanics Meeting.
- [2] Yang Cheng, James Miller "Autonomous landmark Based Spacecraft Navigation", 13th Annual AAS/AIAA Space Flight Mechanics Meeting.
- [3] J. K. Miller, et al., "Determination of Shape, Gravity, and Rotational State of Asteroid 433 Eros", *Icarus* 155 3-17 (2002).
- [4] David Forsyth, et al., "Invariant descriptors for 3-D object recognition and pose" *IEEE PAMI*, Vol. 13, No. 10, 1991.
- [5] Douglas Bernard and Matthew Golombek, "Crater and rockhazard modeling for Mars Landing", AIAA Space 2001 conference, Albuquerque, NM.
- [6] A. Ansar and K. Daniilidis, "Linear Pose Estimation from Points and Lines," *IEEE PAMI*, 25:578-589, 2003.
- [7] B.K.P. Horn, et al., "Closed-form Solution of Absolute Orientation Using Orthonormal Matrices," *Journal Opt. Soc. Am. A*, A5: 1127-1135, 1988.
- [8] C-P. Lu, G. Hager, and E. Mjolsness, "Fast and Globally Convergent Pose Estimation from Video Images," *IEEE PAMI* 22(6): 610-622 (2000)
- [9] Yang Cheng, et al., "The Mars Exploration Rovers Descent Image Motion Estimation System", *IEEE Intelligent Systems*, May/June 2004, 13-21
- [10] W.K. Hartmann, "Martian Cratering, 6, Crater count isochrones and evidence for recent volcanism from Mars Global Surveyor," *Meteoritics and Planetary Science*, 34, 167-177, 1999.